

Staggered Time Average Algorithm for Stochastic Non-smooth Optimization with $O(1/T)$ Convergence

Sucha Supittayapornpong, Michael J. Neely
 Department of Electrical Engineering
 University of Southern California
 Email: supittay@usc.edu, mjneely@usc.edu

Abstract—Stochastic non-smooth convex optimization constitutes a class of problems in machine learning and operations research. This paper considers minimization of a non-smooth function based on stochastic subgradients. When the function has a *locally polyhedral* structure, a staggered time average algorithm is proven to have $O(1/T)$ convergence rate. A more general convergence result is proven when the locally polyhedral assumption is removed. In that case, the convergence bound depends on the curvature of the function near the minimum. Finally, the locally polyhedral assumption is shown to improve convergence beyond $O(1/T)$ for a special case of deterministic problems.

I. INTRODUCTION

Non-smooth convex optimization constitutes a class of problems in machine learning and operations research. Examples include optimization of the hinge loss function [1] (used in support vector machines) and 1-norm regularization in regression problems [2]. A non-smooth function is continuous and non-differentiable [3]. This lack of differentiability makes it challenging to design algorithms with fast convergence.

This paper considers a stochastic optimization problem:

$$\min_{\mathbf{w} \in \mathcal{W}} F(\mathbf{w}) \quad (1)$$

where \mathcal{W} is a closed and convex set and $F : \mathcal{W} \rightarrow \mathbb{R}$ is a continuous, convex (but not necessary strongly convex), and non-smooth function. Function F may not be known. The optimization proceeds by obtaining an unbiased stochastic subgradient of F from an oracle. This model with an oracle has been previously used in literature, such as [4], [5]. Formally, let $\mathbf{g}(\mathbf{w})$ be the subgradient of F at \mathbf{w} . Receiving $\mathbf{w} \in \mathcal{W}$, the oracle gives an unbiased stochastic subgradient $\hat{\mathbf{g}}(\mathbf{w})$ of F at \mathbf{w} satisfying $\mathbb{E}[\hat{\mathbf{g}}(\mathbf{w})|\mathbf{w}] = \mathbf{g}(\mathbf{w})$. Note that $\mathbf{g}(\mathbf{w})$ may not be known, but $\hat{\mathbf{g}}(\mathbf{w})$ is known. An algorithm proceeds by generating a sequence of \mathbf{w}_t vectors that are given to the oracle. The next vector \mathbf{w}_{t+1} is determined as a function of the history of oracle outputs. The history and the $\{\mathbf{w}_t\}$ sequence is used to compute an estimate $\hat{\mathbf{w}}$ of the optimal solution. It has been noted in [5] that this model can be applied to a class of learning problems in [6].

Given $\epsilon > 0$, an estimate $\hat{\mathbf{w}}$ is an $O(\epsilon)$ -approximation if

$$\mathbb{E}[F(\hat{\mathbf{w}})] - \min_{\mathbf{w} \in \mathcal{W}} F(\mathbf{w}) \leq O(\epsilon).$$

Let T be the number of unbiased stochastic subgradients obtained from an oracle. The convergence rate is determined by the rate at which the estimate converges to the true answer, as a function of T . For example, an algorithm with $O(1/\sqrt{T})$ convergence rate provides the estimate whose deviation from optimality decays to zero like:

$$\mathbb{E}[F(\hat{\mathbf{w}})] - \min_{\mathbf{w} \in \mathcal{W}} F(\mathbf{w}) \leq O(1/\sqrt{T}).$$

Prior work in [7], [5] develops algorithms with $O(1/\sqrt{T})$ convergence rate when the function is non-smooth. A smoothing method in [8] improves the convergence rate to $O(1/T)$ when F is a linear combination of a smooth convex function and a non-smooth convex function with special structure. Related improvements can be shown when the non-smooth function is *strongly convex* [4], [5], [9]. The suffix algorithm in [4] is shown to have $O(1/T)$ convergence rate. The work by [5] shows that using the reducing step size leads to $O(\log(T)/T)$ and $O(1/T)$ convergence rates for the last round solution and a solution calculated from a polynomial-decay averaging. Then, the algorithm achieving optimal $O(1/T)$ convergence rate is developed in [9]. Note that all previous results that achieve the $O(1/T)$ convergence rate rely on either a restrictive strong convexity or a special structure of F . It has become an open problem whether the $O(1/T)$ convergence rate can be achieved for a non-smooth and non-strongly convex function.

In this paper, a non-smooth convex function is considered. In the case when the function has a *locally polyhedral* structure¹, a staggered time average algorithm, based on a stochastic subgradient algorithm with constant step size, is proposed. The algorithm calculates the $O(\epsilon)$ -approximation estimates with $O(1/T)$ convergence rate. For a general convex function, the convergence rate depends on the curvature near the minimum. To our knowledge, with the locally polyhedral structure, this is the first $O(1/T)$ convergence rate for a non-smooth convex function, which requires neither strong convexity nor smoothing.

¹The locally polyhedral structure is also called weak sharp minima in previous literature [10]. It is also the generalization of a sharp minimum function in [11]

The paper is organized as follows. Section II provides notations, the staggered time average algorithm, and preliminary results. Sections III and IV prove respectively the results under the locally polyhedral and the general convex structures. Section V shows a fast convergence for deterministic problems. Experiments are performed in Section VI. Section VII concludes the paper.

II. PRELIMINARIES

The closed convex set \mathcal{W} is a subset of \mathbb{R}^N , for some positive integer N , with Euclidean norm $\|\cdot\|$ and inner product $\langle \cdot, \cdot \rangle$. Function F is assumed to be convex (possibly non-smooth) over \mathcal{W} and satisfies the following assumption. Define $F^* \triangleq \inf_{\mathbf{w} \in \mathcal{W}} F(\mathbf{w})$.

Assumption 1. *The minimum of F is achievable in \mathcal{W} , and the set of optimal solutions $\mathcal{W}^* = \text{Arginf}_{\mathbf{w} \in \mathcal{W}} F(\mathbf{w}) = \{\mathbf{w}^* \in \mathcal{W} : F(\mathbf{w}^*) = F^*\}$ is closed.*

A subgradient of F at $\mathbf{w} \in \mathcal{W}$ is denoted by $\mathbf{g}(\mathbf{w})$ and satisfies for any $\mathbf{w}' \in \mathcal{W}$:

$$F(\mathbf{w}') \geq F(\mathbf{w}) + \langle \mathbf{g}(\mathbf{w}), \mathbf{w}' - \mathbf{w} \rangle. \quad (2)$$

An unbiased stochastic subgradient at \mathbf{w} is denoted by $\hat{\mathbf{g}}(\mathbf{w})$, which satisfies $\mathbb{E}[\hat{\mathbf{g}}(\mathbf{w})|\mathbf{w}] = \mathbf{g}(\mathbf{w})$.

Assumption 2. *There exists a constant $G < \infty$ such that*

$$\|\hat{\mathbf{g}}(\mathbf{w})\| \leq G \quad \forall \mathbf{w} \in \mathcal{W}.$$

Assumption 2 is also used in previous literature.

A stochastic subgradient algorithm with a positive constant step size $\alpha > 0$ initializes $\mathbf{w}_0 \in \mathcal{W}$ and proceeds repeatedly as

$$\mathbf{w}_{t+1} = \Pi_{\mathcal{W}}[\mathbf{w}_t - \alpha \hat{\mathbf{g}}(\mathbf{w}_t)] \quad \forall t \in \{0, 1, 2, \dots\}, \quad (3)$$

where $\Pi_{\mathcal{W}}$ denotes projection on \mathcal{W} and \mathbf{w}_t denotes the values of \mathbf{w} at round t .

A. Staggered Time Averages

The staggered time average algorithm is summarized in Algorithm 1. Define:

$$\bar{\mathbf{w}}_{2^k-1}^T = \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{w}_{2^k-1+t}.$$

This average can be computed on-the-fly as shown in Algorithm 1.

Algorithm 1 implements the subgradient algorithm (3) with constant step size in each round. The staggered time averages reset the calculation of estimates every $2^k - 1$ for $k \in \{0, 1, 2, \dots\}$. Specifically, for every $k \in \{0, 1, 2, \dots\}$, the algorithm generates estimates $\bar{\mathbf{w}}_{2^k-1}^T$ for $T \in \{1, \dots, 2^k\}$. To analyze Algorithm 1, the properties of the subgradient algorithm are proven in this section. Then the staggered time averages are analyzed in Section III and Section IV.

Note that Algorithm 1 is different from the suffix averaging in [4] which uses the reducing step size.

Algorithm 1 Staggered Time Averages

```

Initialize:  $\mathbf{w}_0 \in \mathcal{W}$ ,  $\alpha > 0$ 
for  $t \in \{0, 1, 2, \dots\}$  do
  // Staggered time averages
  if  $t = 2^{k^*} - 1$  for some  $k^* \in \{0, 1, 2, \dots\}$  then
     $k \leftarrow k^*$ 
     $\bar{\mathbf{w}}_{2^k-1}^1 \leftarrow \mathbf{w}_t$ 
  else
     $\bar{\mathbf{w}}_{2^k-1}^{t-2^k+2} \leftarrow \frac{t-2^k+1}{t-2^k+2} \bar{\mathbf{w}}_{2^k-1}^{t-2^k+1} + \frac{1}{t-2^k+2} \mathbf{w}_t$ 
  end if
  // Stochastic subgradient
   $\mathbf{w}_{t+1} \leftarrow \Pi_{\mathcal{W}}[\mathbf{w}_t - \alpha \hat{\mathbf{g}}(\mathbf{w}_t)]$ 
end for

```

B. Basic Results

We consider algorithm (3) with a positive constant step size α . The initial $\mathbf{w}_0 \in \mathcal{W}$ is any constant vector. For every $\mathbf{w}_t \in \mathcal{W}$, define the closest optimal solution to \mathbf{w}_t as

$$\mathbf{w}_t^* = \text{arginf}_{\mathbf{w} \in \mathcal{W}^*} \|\mathbf{w} - \mathbf{w}_t\|.$$

Under Assumption 1, this \mathbf{w}_t^* is unique because of the convexity and the closeness of \mathcal{W}^* . The following lemma modifies a well known manipulation.

Lemma 1. *Suppose Assumptions 1 and 2 hold. It holds for any $t \in \{0, 1, 2, \dots\}$ that*

$$\mathbb{E} \left[\|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\|^2 \middle| \mathbf{w}_t \right] \leq \|\mathbf{w}_t - \mathbf{w}_t^*\|^2 + \alpha^2 G^2 + 2\alpha[F^* - F(\mathbf{w}_t)].$$

Proof. For any t , by definition of \mathbf{w}_{t+1}^* as the minimizer of $\|\mathbf{w} - \mathbf{w}_{t+1}\|^2$ over all $\mathbf{w} \in \mathcal{W}$, we have:

$$\begin{aligned} \|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\|^2 &\leq \|\mathbf{w}_{t+1} - \mathbf{w}_t^*\|^2 \\ &= \|\Pi_{\mathcal{W}}[\mathbf{w}_t - \alpha \hat{\mathbf{g}}(\mathbf{w}_t)] - \mathbf{w}_t^*\|^2 \\ &\leq \|\mathbf{w}_t - \alpha \hat{\mathbf{g}}(\mathbf{w}_t) - \mathbf{w}_t^*\|^2 \end{aligned}$$

where the final inequality holds by the non-expansive property of the projection onto the convex set \mathcal{W} [12]. Expanding the right-hand-side gives:

$$\begin{aligned} \|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\|^2 &= \|\mathbf{w}_t - \mathbf{w}_t^*\|^2 + \alpha^2 \|\hat{\mathbf{g}}(\mathbf{w}_t)\|^2 - 2\alpha \langle \hat{\mathbf{g}}(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}_t^* \rangle \\ &\leq \|\mathbf{w}_t - \mathbf{w}_t^*\|^2 + \alpha^2 G^2 + 2\alpha \langle \hat{\mathbf{g}}(\mathbf{w}_t), \mathbf{w}_t^* - \mathbf{w}_t \rangle \end{aligned}$$

Taking a conditional expectation given \mathbf{w}_t yields

$$\mathbb{E} \left[\|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\|^2 \middle| \mathbf{w}_t \right] \leq \|\mathbf{w}_t - \mathbf{w}_t^*\|^2 + \alpha^2 G^2 + 2\alpha \langle \mathbf{g}(\mathbf{w}_t), \mathbf{w}_t^* - \mathbf{w}_t \rangle.$$

Using the subgradient property in (2) and the fact that $F(\mathbf{w}_t^*) = F^*$ proves the lemma. \square

Note that Lemma 1 uses a projection technique similar to standard analysis for the subgradient projection algorithm, (as in [13], [3] or [12]). The standard approach compares

the current iterate to a fixed optimal point \mathbf{w}^* . Lemma 1 compares to the closest point in the optimal set. This is a simple but important distinction that is crucial in later sections for improved convergence time results.

While Lemma 1 is stated in a form useful for the analysis of later sections, it can readily be used to establish the standard $O(1/\epsilon^2)$ result for convex functions (see also [13], [3], [12]).

Define an average of T -consecutive solutions from t_0 as

$$\bar{\mathbf{w}}_{t_0}^T \triangleq \frac{1}{T} \sum_{t=t_0}^{t_0+T-1} \mathbf{w}_t. \quad (4)$$

Taking an expectation of the result in Lemma 1 gives

$$\mathbb{E}[F(\mathbf{w}_t)] - F^* \leq \frac{\alpha G^2}{2} + \frac{1}{2\alpha} \mathbb{E}[\|\mathbf{w}_t - \mathbf{w}_t^*\|^2 - \|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\|^2].$$

Summing from t_0 to $t_0 + T - 1$ and dividing by T gives

$$\begin{aligned} & \frac{1}{T} \sum_{t=t_0}^{t_0+T-1} \mathbb{E}[F(\mathbf{w}_t)] - F^* \leq \\ & \frac{\alpha G^2}{2} + \frac{1}{2\alpha T} \mathbb{E}[\|\mathbf{w}_{t_0} - \mathbf{w}_{t_0}^*\|^2 - \|\mathbf{w}_{t_0+T} - \mathbf{w}_{t_0+T}^*\|^2]. \end{aligned}$$

Using Jensen's inequality and convexity of F , definition (4), and non-negativity of $\|\mathbf{w}_{t_0+T} - \mathbf{w}_{t_0+T}^*\|^2$ yield:

$$\mathbb{E}[F(\bar{\mathbf{w}}_{t_0}^T)] - F^* \leq \frac{\alpha G^2}{2} + \frac{1}{2\alpha T} \mathbb{E}[\|\mathbf{w}_{t_0} - \mathbf{w}_{t_0}^*\|^2], \quad (5)$$

for any $t_0 \in \{0, 1, 2, \dots\}$ and any positive integer T .

Equation (5) suggests that, to achieve an $O(\epsilon)$ -approximation estimate, one can choose step size $\alpha = \Theta(\epsilon)$, number of rounds $T = \Theta(1/\epsilon^2)$, and define $\bar{\mathbf{w}}$ as the average of \mathbf{w}_t values over the first T rounds. This is equivalent to $O(1/\sqrt{T})$ convergence rate. Fortunately, this convergence rate can be improved by starting the average at an appropriate time depending on the structure of the function F . These structures are shown in Section III and Section IV. We first prove several useful results used in those sections.

C. Concentration Bound

These results are used to upper bound $\mathbb{E}[\|\mathbf{w}_{t_0} - \mathbf{w}_{t_0}^*\|^2]$ in (5). Define $K_t \triangleq \|\mathbf{w}_t - \mathbf{w}_t^*\|$ for all $t \in \{0, 1, 2, \dots\}$.

Lemma 2. Suppose Assumption 2 holds. It holds for every $t \in \{0, 1, 2, \dots\}$ that

$$|K_{t+1} - K_t| \leq 2\alpha G \quad (6)$$

$$\mathbb{E}[K_{t+1} - K_t | \mathbf{w}_t] \leq 2\alpha G. \quad (7)$$

Proof. The first part is proven in two cases.

i) If $K_{t+1} \geq K_t$, definition of \mathbf{w}_{t+1} in (3) and the non-expansive projection implies

$$\begin{aligned} |K_{t+1} - K_t| &= K_{t+1} - K_t \leq \|\mathbf{w}_{t+1} - \mathbf{w}_t^*\| - K_t \\ &\leq \|\mathbf{w}_t - \alpha \hat{\mathbf{g}}(\mathbf{w}_t) - \mathbf{w}_t^*\| - K_t \\ &\leq \|\mathbf{w}_t - \mathbf{w}_t^*\| + \alpha \|\hat{\mathbf{g}}(\mathbf{w}_t)\| - K_t \leq \alpha G. \end{aligned}$$




Fig. 1. Structures of function F

ii) If $K_{t+1} < K_t$, we have

$$\begin{aligned} |K_{t+1} - K_t| &= K_t - K_{t+1} \\ &= \|\mathbf{w}_t - \mathbf{w}_{t+1} + \mathbf{w}_{t+1} - \mathbf{w}_{t+1}^* + \mathbf{w}_{t+1}^* - \mathbf{w}_t^*\| - K_{t+1} \\ &\leq \|\mathbf{w}_t - \mathbf{w}_{t+1}\| + K_{t+1} + \|\mathbf{w}_{t+1}^* - \mathbf{w}_t^*\| - K_{t+1} \\ &\leq 2\|\mathbf{w}_t - \mathbf{w}_{t+1}\| \leq 2\alpha \|\hat{\mathbf{g}}(\mathbf{w}_t)\| \leq 2\alpha G, \end{aligned}$$

where the last line uses non-expansive projection implying $\|\mathbf{w}_{t+1}^* - \mathbf{w}_t^*\| \leq \|\mathbf{w}_{t+1} - \mathbf{w}_t\|$.

These two cases prove the first part. The second part follows by taking a conditional expectation given \mathbf{w}_t of $K_{t+1} - K_t \leq |K_{t+1} - K_t|$. \square

The concentration bound reinterprets the lemma in [14].

Lemma 3. Suppose Assumption 2 holds and there exists a positive real-valued β and a $\gamma \in \mathbb{R}$ such that for all $t \in \{0, 1, 2, \dots\}$:

$$\mathbb{E}[K_{t+1} - K_t | K_t] \leq \begin{cases} 2\alpha G & , \text{if } K_t < \gamma \\ -\beta & , \text{if } K_t \geq \gamma. \end{cases} \quad (8)$$

Assume $K_{t_0} = k_{t_0}$ (with probability 1) for some $k_{t_0} \in \mathbb{R}$. Then following holds for all $t \in \{t_0, t_0 + 1, t_0 + 2, \dots\}$

$$\mathbb{E}[e^{rK_t}] \leq D + (e^{rk_{t_0}} - D)\rho^{t-t_0}, \quad (9)$$

where constants r, ρ , and D are:

$$\begin{aligned} r &\triangleq \frac{3\beta}{12\alpha^2 G^2 + 2\alpha G\beta}, & \rho &\triangleq 1 - \frac{r\beta}{2}, \\ D &\triangleq \frac{(e^{2\alpha G r} - \rho)e^{r\gamma}}{1 - \rho}. \end{aligned}$$

It can be shown that (6) and (8) together imply $0 < \beta \leq 2\alpha G$, and hence it can be shown that $0 < \rho < 1$.

Constants β and γ in Lemma 3 depend on a structure of function F . We then look at the first structure.

III. LOCALLY POLYHEDRAL STRUCTURE

In this section, function F is assumed to have a *locally polyhedral* structure, which is illustrated in Figure 1. This structure is generalized from [15]. It is assumed throughout that Assumptions 1 and 2 still hold. Note that, in machine learning, this F can be the hinge loss function with 1-norm regularization.

Assumption 3. (Locally polyhedral assumption) There exists a constant $L_P > 0$ such that for every $\mathbf{w}_t \in \mathcal{W}$ the following holds

$$F(\mathbf{w}_t) - F^* \geq L_P \|\mathbf{w}_t - \mathbf{w}_t^*\|. \quad (10)$$

The subscript ‘‘P’’ in L_P represents ‘‘polyhedral.’’

A. Drift and Transient Time

Using this locally polyhedral structure, sequence $\{\mathbf{w}_t\}_{t=0}^\infty$ generated by algorithm (3) has the following drift property. Define

$$B_P \triangleq \max \left[\frac{L_P}{2}, \frac{G^2}{L_P} \right].$$

Lemma 4. Suppose Assumptions 1, 2, 3 hold. For any $\mathbf{w}_t \in \mathcal{W}$ and $\|\mathbf{w}_t - \mathbf{w}_t^*\| \geq \alpha B_P$, the following holds

$$\mathbb{E} [\|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\| | \mathbf{w}_t] \leq \|\mathbf{w}_t - \mathbf{w}_t^*\| - \frac{\alpha L_P}{2}. \quad (11)$$

Proof. For any $\mathbf{w}_t \in \mathcal{W}$, if condition

$$\alpha^2 G^2 + 2\alpha[F^* - F(\mathbf{w}_t)] \leq -\alpha L_P \|\mathbf{w}_t - \mathbf{w}_t^*\| + \frac{\alpha^2 L_P^2}{4} \quad (12)$$

is true, then the result in Lemma 1 implies that

$$\begin{aligned} \mathbb{E} [\|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\|^2 | \mathbf{w}_t] \\ \leq \|\mathbf{w}_t - \mathbf{w}_t^*\|^2 - \alpha L_P \|\mathbf{w}_t - \mathbf{w}_t^*\| + \frac{\alpha^2 L_P^2}{4} \\ = \left(\|\mathbf{w}_t - \mathbf{w}_t^*\| - \frac{\alpha L_P}{2} \right)^2. \end{aligned}$$

Applying Jensen's inequality on the left-hand-side gives

$$(\mathbb{E} [\|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\| | \mathbf{w}_t])^2 \leq \left(\|\mathbf{w}_t - \mathbf{w}_t^*\| - \frac{\alpha L_P}{2} \right)^2.$$

When $\|\mathbf{w}_t - \mathbf{w}_t^*\| \geq \alpha B_P \geq \alpha L_P/2$, inequality (11) holds.

It remains to show that condition (12) must hold whenever $\|\mathbf{w}_t - \mathbf{w}_t^*\| \geq \alpha B_P$. Starting with the left-hand-side of (12) we have, for every $\mathbf{w}_t \in \mathcal{W}$:

$$\alpha^2 G^2 + 2\alpha[F^* - F(\mathbf{w}_t)] \leq \alpha^2 G^2 - 2\alpha L_P \|\mathbf{w}_t - \mathbf{w}_t^*\|$$

where the inequality follows the locally polyhedral structure (10). By the definition of B_P , it holds that $\|\mathbf{w}_t - \mathbf{w}_t^*\| \geq \alpha G^2/L_P$ and $\mathbf{w}_t \in \mathcal{W}$. Substituting this into the above gives the result. \square

This lemma implies that, when the distance between \mathbf{w}_t and \mathbf{w}_t^* is at least αB_P , then \mathbf{w}_{t+1} is expected to get closer to \mathbf{w}_{t+1}^* . This phenomenon suggests that \mathbf{w}_t will concentrate around \mathcal{W}^* after some transient time, if it is not inside the set already.

Let constants U_P, r_P, ρ_P, D_P be

$$U_P \triangleq \frac{2(D_P + 1)}{r_P^2}, \quad (13)$$

$$r_P \triangleq \frac{3L_P}{24G^2 + 2L_P G}, \quad (14)$$

$$\rho_P \triangleq 1 - \frac{3L_P^2}{4(24G^2 + 2L_P G)}, \quad (15)$$

$$D_P \triangleq \frac{(e^{2Gr_P} - \rho_P)e^{r_P B_P}}{1 - \rho_P}. \quad (16)$$

Given any $\mathbf{w}_0 \in \mathcal{W}$, define

$$T_P(\alpha) \triangleq \left\lceil \frac{r_P \|\mathbf{w}_0 - \mathbf{w}_0^*\|}{\alpha \log(1/\rho_P)} \right\rceil, \quad (17)$$

where constants r_P and ρ_P are defined in (14) and (15). This $T_P(\alpha)$ can be called a transient time for the locally polyhedral structure, since a useful bound holds after this time.

Lemma 5. Suppose Assumptions 1, 2, 3 hold. When $t \geq T_P(\alpha)$, the following holds

$$\mathbb{E} [\|\mathbf{w}_t - \mathbf{w}_t^*\|^2] \leq \alpha^2 U_P,$$

where constant U_P is defined in (13).

Proof. Lemma 4 gives $\mathbb{E}[K_{t+1} - K_t | K_t] \leq -\alpha L_P/2$ as in (11) when $K_t = \|\mathbf{w}_t - \mathbf{w}_t^*\| \geq \alpha B_P$. Therefore, the constants β and γ in Lemma 3 can be set as $\beta = \alpha L_P/2$ and $\gamma = \alpha B_P$. When $t_0 = 0$, we have $k_{t_0} = K_0 = \|\mathbf{w}_0 - \mathbf{w}_0^*\|$ (with probability 1). From (9), it holds for all $t \in \{0, 1, 2, \dots\}$ that

$$\begin{aligned} \mathbb{E} \left[e^{\frac{r_P K_t}{\alpha}} \right] &\leq D_P + \left(e^{\frac{r_P k_{t_0}}{\alpha}} - D_P \right) \rho_P^{(t-0)} \\ &\leq D_P + e^{\frac{r_P K_0}{\alpha}} \rho_P^t, \end{aligned} \quad (18)$$

where constants r_P, ρ_P, D_P are defined in (14)–(16) respectively. We then show that

$$e^{\frac{r_P K_0}{\alpha}} \rho_P^t \leq 1 \quad \forall t \geq T_P(\alpha). \quad (19)$$

Inequality $e^{\frac{r_P K_0}{\alpha}} \rho_P^t \leq 1$ is equivalent to $t \geq \frac{r_P K_0}{\alpha \log(1/\rho_P)}$ by arithmetic and the fact that $\log(1/\rho_P) > 0$. From the definition of $T_P(\alpha)$ in (17), it holds that $T_P(\alpha) \geq \frac{r_P K_0}{\alpha \log(1/\rho_P)}$, and the results (19) follows.

From (19), inequality (18) becomes

$$\mathbb{E} \left[e^{\frac{r_P K_t}{\alpha}} \right] \leq D_P + 1 \quad \forall t \geq T_P(\alpha).$$

The Chernoff bounds (see, for example, [16]) implies for any $m > 0$ that

$$\begin{aligned} \mathbb{P}\{K_t \geq m\} &\leq e^{-\frac{r_P m}{\alpha}} \mathbb{E} \left[e^{\frac{r_P K_t}{\alpha}} \right] \\ &\leq e^{-\frac{r_P m}{\alpha}} (D_P + 1) \quad \forall t \geq T_P(\alpha). \end{aligned}$$

Using $\mathbb{E}[K_t^2] = 2 \int_0^\infty m \mathbb{P}\{K_t \geq m\} dm$ and the above bound proves the lemma. \square

The definition of $T_P(\alpha)$ in (17) implies that the transient time is $O(1/\alpha)$ under the locally polyhedral structure. Then, Lemma 5 implies that $\mathbb{E} [\|\mathbf{w}_t - \mathbf{w}_t^*\|^2] \leq \alpha^2 U_P$ every round t after the transient time.

B. Convergence Rate

We are now ready to prove the convergence rate of the staggered time averages in Algorithm 1 under the locally polyhedral structure.

Theorem 1. Suppose Assumptions 1, 2, 3 hold. It holds for any $t \geq T_P(\alpha)$ and any positive integer T that

$$\mathbb{E} [F(\bar{\mathbf{w}}_t^T)] - F^* \leq \frac{\alpha G^2}{2} + \frac{\alpha U_P}{2T}, \quad (20)$$

where constant U_P is defined in (13).

Proof. The theorem follows from (5) and Lemma 5 that

$$\mathbb{E} [F(\bar{\mathbf{w}}_t^T)] - F^* \leq \frac{\alpha G^2}{2} + \frac{\alpha^2 U_P}{2\alpha T} \quad \forall t \geq T_P(\alpha). \quad \square$$

After the transient time, Theorem 1 implies that estimates, as the averages, converge as $O(\alpha + \alpha/T)$. To obtain an $O(\epsilon)$ -approximation estimate, the step size must be set to $\Theta(\epsilon)$. Recall that the transient time (17) is $O(1/\alpha)$ and that Algorithm 1 resets the averages at round $2^k - 1$ for $k \in \{0, 1, 2, \dots\}$. Let $\hat{k} = \arg\inf_{k \in \{0, 1, 2, \dots\}} 2^k - 1 \geq T_P(\alpha)$ be the first reset after the transient time. The exponential increasing implies that $2^{\hat{k}} - 1 \leq 2T_P(\alpha)$. Therefore, the total time to obtain the estimate is $O(1/\epsilon)$, and the convergence rate is $O(1/T)$. Note that the staggered time average algorithm is proposed, because $\|\mathbf{w}_0 - \mathbf{w}_0^*\|$ in (17) can not be upper bounded if \mathcal{W} is unbounded. Also, even though the convergence rate depends only on the step size, performing the averages also helps obtaining more accurate estimates, as shown in Section VI.

IV. GENERAL CONVEX FUNCTION

In this section, function F is allowed to be a *general convex* function, possibly one that does not satisfy the locally polyhedral structure (Assumption 3). A general convex function is illustrated in Figure 1. It is assumed throughout that Assumptions 1 and 2 still hold.

Define $\mathcal{A}(S) = \{\mathbf{w}_t \in \mathcal{W} : \|\mathbf{w}_t - \mathbf{w}_t^*\| = S\}$, and define S_{\max} as the supremum value of $S > 0$ for which $\mathcal{A}(S)$ is nonempty (S_{\max} is possibly infinity). Assume that $S_{\max} > 0$. Convexity of the set \mathcal{W} implies that $\mathcal{A}(S)$ is nonempty for all $S \in (0, S_{\max})$. For each $S \in (0, S_{\max})$ define:

$$\eta(S) = \inf_{\mathbf{w}_t \in \mathcal{W} : \|\mathbf{w}_t - \mathbf{w}_t^*\| = S} \frac{|F(\mathbf{w}_t) - F^*|}{\|\mathbf{w}_t - \mathbf{w}_t^*\|}. \quad (21)$$

Lemma 6. *Suppose Assumption 1 holds. If F is convex and \mathcal{W} is closed, then for all $S \in (0, S_{\max})$:*

i) $\eta(S) > 0$

ii) *Whenever $\mathbf{w}_t \in \mathcal{W}$ and $\|\mathbf{w}_t - \mathbf{w}_t^*\| \geq S$, it holds that*

$$F(\mathbf{w}_t) - F^* \geq \eta(S) \|\mathbf{w}_t - \mathbf{w}_t^*\|. \quad (22)$$

Proof. The first part is proven by contradiction. Define $\mathcal{A} = \{\mathbf{w}_t \in \mathcal{W} : \|\mathbf{w}_t - \mathbf{w}_t^*\| = S\}$, which is nonempty. Note that \mathcal{A} is compact. Suppose $\eta(S) = 0$. Function $|F(\mathbf{w}_t) - F^*|/S$ is continuous. The infimum of a continuous function over the compact set is achieved by a point in the set. Thus, there is a point $\mathbf{y} \in \mathcal{A}$ such that $|F(\mathbf{y}) - F^*|/S = 0$. That is $F(\mathbf{y}) = F^*$, and $\mathbf{y} \in \mathcal{W}^*$. Since $\mathbf{y} \in \mathcal{A}$, it also satisfies $\inf_{\mathbf{y}^* \in \mathcal{W}^*} \|\mathbf{y} - \mathbf{y}^*\| = S$ and $\mathbf{y} \notin \mathcal{W}^*$, which is a contradiction.

The second part is proven using the convexity of F . Let $\mathbf{z} \in \mathcal{W}$ be a vector such that $\|\mathbf{z} - \mathbf{z}^*\| \geq S$ where $\mathbf{z}^* = \arg\inf_{\hat{\mathbf{z}} \in \mathcal{W}^*} \|\mathbf{z} - \hat{\mathbf{z}}\|$. We want to show that $F(\mathbf{z}) - F^* \geq \eta(S) \|\mathbf{z} - \mathbf{z}^*\|$. The convexity of the set \mathcal{W} implies that the line segment between \mathbf{z} and \mathbf{z}^* is inside \mathcal{W} . The convexity of F over \mathcal{W} implies that F is convex when it is restricted to this line segment.

Define \mathbf{y} as a point on this line segment such that $\|\mathbf{y} - \mathbf{y}^*\| = S$ where $\mathbf{y}^* = \arg\inf_{\hat{\mathbf{y}} \in \mathcal{W}^*} \|\mathbf{y} - \hat{\mathbf{y}}\|$. Then both $\mathbf{y} \in \mathcal{W}$ and $\mathbf{y} \in \mathcal{A}$. The convexity of F over the line segment implies that

$$\frac{F(\mathbf{z}) - F^*}{\|\mathbf{z} - \mathbf{z}^*\|} \geq \frac{F(\mathbf{y}) - F^*}{\|\mathbf{y} - \mathbf{y}^*\|} \geq \eta(S),$$

where the last inequality uses (21). \square

The difference between Assumption 3 and Lemma 6 is that the bound (22) only holds when $\|\mathbf{w}_t - \mathbf{w}_t^*\| \geq S$. The choice of S for a particular function F affects the transient time and convergence of achieving $O(\epsilon)$ -approximation estimates. This effect does not occur with Assumption 3.

A. Drift and Transient Time

Using Lemma 6, the sequence $\{\mathbf{w}_t\}_{t=0}^\infty$ generated by algorithm (3) has the following property. Define

$$B_G(\alpha, S) \triangleq \max \left[\frac{\alpha \eta(S)}{2}, S, \frac{\alpha G^2}{\eta(S)} \right]. \quad (23)$$

Lemma 7. *Suppose Assumptions 1 and 2 hold. Function F is convex with $\eta(S)$ defined in (21) for all $S \in (0, S_{\max})$. For any $\mathbf{w}_t \in \mathcal{W}$ that $\|\mathbf{w}_t - \mathbf{w}_t^*\| \geq B_G(\alpha, S)$, the following holds*

$$\mathbb{E} [\|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\| | \mathbf{w}_t] \leq \|\mathbf{w}_t - \mathbf{w}_t^*\| - \frac{\alpha \eta(S)}{2}. \quad (24)$$

Proof. For any $\mathbf{w}_t \in \mathcal{W}$, if condition

$$\alpha^2 G^2 + 2\alpha[F^* - F(\mathbf{w}_t)] \leq -\alpha \eta(S) \|\mathbf{w}_t - \mathbf{w}_t^*\| + \frac{\alpha^2 \eta(S)^2}{4} \quad (25)$$

is true, then the result in Lemma 1 implies that

$$\begin{aligned} & \mathbb{E} [\|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\|^2 | \mathbf{w}_t] \\ & \leq \|\mathbf{w}_t - \mathbf{w}_t^*\|^2 - \alpha \eta(S) \|\mathbf{w}_t - \mathbf{w}_t^*\| + \frac{\alpha^2 \eta(S)^2}{4} \\ & = \left[\|\mathbf{w}_t - \mathbf{w}_t^*\| - \frac{\alpha \eta(S)}{2} \right]^2. \end{aligned}$$

Applying Jensen's inequality on the left-hand-side gives

$$(\mathbb{E} [\|\mathbf{w}_{t+1} - \mathbf{w}_{t+1}^*\| | \mathbf{w}_t])^2 \leq \left[\|\mathbf{w}_t - \mathbf{w}_t^*\| - \frac{\alpha \eta(S)}{2} \right]^2.$$

When $\|\mathbf{w}_t - \mathbf{w}_t^*\| \geq B_G(\alpha, S) \geq \alpha \eta(S)/2$, inequality (24) holds.

It remains to show condition (25) must hold whenever $\|\mathbf{w}_t - \mathbf{w}_t^*\| \geq B_G(\alpha, S)$. Starting with the left-hand-side of (25), we have that

$$\begin{aligned} \alpha^2 G^2 + 2\alpha[F^* - F(\mathbf{w}_t)] & \leq \\ & \alpha^2 G^2 - 2\alpha \eta(S) \|\mathbf{w}_t - \mathbf{w}_t^*\| \end{aligned}$$

where the inequality follows (22), as $B_G(\alpha, S) \geq S$. By the definition of $B_G(\alpha, S)$, it holds that $\|\mathbf{w}_t - \mathbf{w}_t^*\| \geq \alpha G^2 / \eta(S)$. Substituting this into the above inequality gives the result. \square

This result is similar to Lemma 4 except that $B_G(\alpha, S)$ depends on both α and S unlike αB_P in the locally polyhedral case.

Let constants $U_G(\alpha, S)$, $r_G(S)$, $\rho_G(S)$, $D_G(\alpha, S)$ be

$$U_G(\alpha, S) \triangleq \frac{2[D_G(\alpha, S) + 1]}{r_G(S)^2}, \quad (26)$$

$$r_G(S) \triangleq \frac{3\eta(S)}{24G^2 + 2\eta(S)G}, \quad (27)$$

$$\rho_G(S) \triangleq 1 - \frac{3\eta(S)^2}{4[24G^2 + 2\eta(S)G]}, \quad (28)$$

$$D_G(\alpha, S) \triangleq \frac{[e^{2Gr_G(S)} - \rho_G(S)] e^{\frac{r_G(S)B_G(\alpha, S)}{\alpha}}}{1 - \rho_G(S)}. \quad (29)$$

Define the transient time for a general convex function as

$$T_G(\alpha, S) \triangleq \left\lceil \frac{r_G(S)\|\mathbf{w}_0 - \mathbf{w}_0^*\|}{\alpha \log(1/\rho_G(S))} \right\rceil, \quad (30)$$

where $r_G(S)$ and $\rho_G(S)$ are defined in (27) and (28).

Lemma 8. Suppose Assumptions 1 and 2 hold. Function F is convex with $\eta(S)$ defined in (21) for all $S \in (0, S_{\max})$. When $t \geq T_G(\alpha, S)$, the following holds

$$\mathbb{E} [\|\mathbf{w}_t - \mathbf{w}_t^*\|^2] \leq \alpha^2 U_G(\alpha, S),$$

where constant $U_G(\alpha, S)$ is defined in (26).

Proof. From Lemma 3, the constants are $\beta = \alpha\eta(S)/2$ and $\gamma = B_G(\alpha, S)$, where (8) holds due to Lemma 7. When $t_0 = 0$, we have $k_{t_0} = K_0 = \|\mathbf{w}_0 - \mathbf{w}_0^*\|$ (with probability 1). From (9), it holds for all $t \in \{0, 1, 2, \dots\}$ that

$$\begin{aligned} \mathbb{E} \left[e^{\frac{r_G(S)K_t}{\alpha}} \right] &\leq D_G(\alpha, S) \\ &\quad + \left[e^{\frac{r_G(S)k_{t_0}}{\alpha}} - D_G(\alpha, S) \right] \rho_G(S)^{(t-0)} \\ &\leq D_G(\alpha, S) + e^{\frac{r_G(S)K_0}{\alpha}} \rho_G(S)^t, \end{aligned} \quad (31)$$

where constants $r_G(S)$, $\rho_G(S)$, $D_G(\alpha, S)$ are defined in (27)–(29) respectively. We then show that

$$e^{\frac{r_G(S)K_0}{\alpha}} \rho_G(S)^t \leq 1 \quad \forall t \geq T_G(\alpha, S). \quad (32)$$

Inequality $e^{\frac{r_G(S)K_0}{\alpha}} \rho_G(S)^t \leq 1$ is equivalent to $t \geq \frac{r_G(S)K_0}{\alpha \log(1/\rho_G(S))}$. From the definition of $T_G(\alpha, S)$ in (30), it holds that $T_G(\alpha, S) \geq \frac{r_G(S)K_0}{\alpha \log(1/\rho_G(S))}$, and the results (32) follows.

From (32), inequality (31) becomes

$$\mathbb{E} \left[e^{\frac{r_G(S)K_t}{\alpha}} \right] \leq D_G(\alpha, S) + 1 \quad \forall t \geq T_G(\alpha, S).$$

The Chernoff bounds implies for any $m > 0$ that

$$\mathbb{P} \{K_t \geq m\} \leq e^{-\frac{r_G(S)m}{\alpha}} [D_G(\alpha, S) + 1] \quad \forall t \geq T_G(\alpha, S).$$

Using $\mathbb{E} [K_t^2] = 2 \int_0^\infty m \mathbb{P} \{K_t \geq m\} dm$ and the above bound proves the lemma. \square

The definition of $T_G(\alpha, S)$ in (30) implies that the transient time for a general convex function depends on a step size

and the curvature near the unique minimum. Then, Lemma 8 implies that $\mathbb{E} [\|\mathbf{w}_t - \mathbf{w}_t^*\|^2] \leq \alpha^2 U_G(\alpha, S)$ every round t after the transient time.

B. Convergence Rate

We are now ready to prove the convergence rate of the staggered time averages in Algorithm 1 under a general convex function.

Theorem 2. Suppose Assumptions 1 and 2 hold. Function F is convex with $\eta(S)$ defined in (21) for all $S \in (0, S_{\max})$. It holds for any $t \geq T_G(\alpha, S)$ and any positive integer T that

$$\mathbb{E} [F(\bar{\mathbf{w}}_t^T)] - F^* \leq \frac{\alpha G^2}{2} + \frac{\alpha U_G(\alpha, S)}{2T}, \quad (33)$$

where constant $U_G(\alpha, S)$ is defined in (26)

Proof. The theorem follows from (5) and Lemma 8 that

$$\mathbb{E} [F(\bar{\mathbf{w}}_t^T)] - F^* \leq \frac{\alpha G^2}{2} + \frac{\alpha^2 U_G(\alpha, S)}{2\alpha T} \quad \forall t \geq T_G(\alpha, S). \quad \square$$

The transient time (30) and Theorem 2 can be interpreted as a class of convergence bounds that can be optimized over any $S \in (0, S_{\max})$. Indeed, the values of S near the minimum of F plays a crucial role in (23), which affects much of the analysis in this section.

V. FAST CONVERGENCE FOR DETERMINISTIC PROBLEMS

This section revisits problems with the locally polyhedral structure, so that Assumptions 1, 2, 3 hold. However, it considers a *deterministic* scenario where the oracle returns the exact subgradient, rather than an unbiased stochastic subgradient. It is shown that a variation on the basic algorithm that uses a variable step size can significantly improve the convergence rate. Specifically, fix $\epsilon > 0$. The basic algorithm of Section III produces an $O(\epsilon)$ -approximation within $O(1/\epsilon)$ rounds. The modified algorithm of this section does the same $O(\epsilon)$ -approximation with only $O(\log(1/\epsilon))$ rounds. In particular, this is faster than the lower bound $\Omega(1/\sqrt{T})$ for a non-smooth function with Lipschitz continuity in [3]. This does not contradict the Nesterov result in [3], because that result shows the existence of a function with $\Omega(1/\sqrt{T})$ convergence rate, while the locally polyhedral structure does not fall into a class of that function. Interestingly, the algorithm in this section is faster than other algorithms with $O(1/T^2)$ convergence rates [3], [17].

Assume the function $F(\mathbf{w})$ is Lipschitz continuous over $\mathbf{w} \in \mathcal{W}$ with Lipschitz constant $H > 0$, so that:

$$|F(\mathbf{w}) - F(\mathbf{w}')| \leq H \|\mathbf{w} - \mathbf{w}'\| \quad \forall \mathbf{w}, \mathbf{w}' \in \mathcal{W}. \quad (34)$$

Assume there is a known positive value $Z < \infty$ such that $\|\mathbf{w}_0 - \mathbf{w}_0^*\| \leq Z$. Fix $\epsilon > 0$, and fix M as any positive integer. The idea is to run the algorithm over successive frames. Label the frames $i \in \{1, 2, 3, \dots, M\}$. Let $\mathbf{w}_{[i]}$

be the initial vector in \mathcal{W} at the start of frame i . Define $\mathbf{w}_{[1]} = \mathbf{w}_0$. Define constants U_P, r_P, ρ_P, D_P as in (13)–(16). Define $\theta = \max[\sqrt{U_P}, Z]$, and define the frame size T as:

$$T = \left\lceil \frac{2r_P\theta}{\log(1/\rho_P)} \right\rceil$$

The algorithm for each frame $i \in \{1, 2, 3, \dots\}$ is:

- Define the step size for frame i as $\alpha_{[i]} = 2^{-i}$.
- Run the constant step size algorithm (3) using step size $\alpha_{[i]}$ over T rounds, using initial vector $\mathbf{w}_{[i]}$.
- Define $\mathbf{w}_{[i+1]}$ as the \mathbf{w}_t vector computed in the last round of frame i .

Notice that the completion of M frames requires $MT = O(M)$ rounds. The vector computed in the last round of the last frame is defined as $\mathbf{w}_{[M+1]}$. The next theorem shows that this vector is indeed an $O(2^{-M})$ -approximation.

Theorem 3. *In the deterministic setting and when Assumptions 1, 2, 3 hold, the final vector $\mathbf{w}_{[M+1]}$ satisfies:*

$$\|\mathbf{w}_{[M+1]} - \mathbf{w}_{[M+1]}^*\| \leq \theta 2^{-M} \quad (35)$$

$$F(\mathbf{w}_{[M+1]}) - F^* \leq \theta H 2^{-M} \quad (36)$$

Proof. The proof is by induction on the rounds $i \in \{1, 2, \dots, M\}$. Assume the following holds on a given $i \in \{1, 2, \dots, M\}$:

$$\|\mathbf{w}_{[i]} - \mathbf{w}_{[i]}^*\| \leq \theta 2^{-(i-1)} \quad (37)$$

This holds by assumption on the first frame $i = 1$. We now show (37) holds for $i + 1$. The goal is to use Lemma 5 with initial condition $\mathbf{w}_{[i]}$. Since the step size is $\alpha_{[i]}$ for this frame, the value $T_P(\alpha_{[i]})$ defined in (17) satisfies:

$$T_P(\alpha_{[i]}) = \left\lceil \frac{r_P \|\mathbf{w}_{[i]} - \mathbf{w}_{[i]}^*\|}{\alpha_{[i]} \log(1/\rho_P)} \right\rceil \leq \left\lceil \frac{r_P \theta 2^{-(i-1)}}{\alpha_{[i]} \log(1/\rho_P)} \right\rceil = T$$

where the inequality holds by the induction assumption (37), and the last equality holds by definition of $\alpha_{[i]}$. Recall that $\mathbf{w}_{[i+1]}$ is defined as the final \mathbf{w}_t value after T rounds of the frame. It follows by Lemma 5 that:

$$\mathbb{E} \left[\|\mathbf{w}_{[i+1]} - \mathbf{w}_{[i+1]}^*\|^2 \right] \leq \alpha_{[i]}^2 U_P \leq \alpha_{[i]}^2 \theta^2$$

On the other hand, this deterministic setting produces a deterministic sequence, so that all expectations can be removed:

$$\|\mathbf{w}_{[i+1]} - \mathbf{w}_{[i+1]}^*\|^2 \leq \alpha_{[i]}^2 \theta^2$$

Taking a square root and using the definition of $\alpha_{[i]}$ proves:

$$\|\mathbf{w}_{[i+1]} - \mathbf{w}_{[i+1]}^*\| \leq \theta 2^{-i}$$

This completes the induction, so that (37) holds for all $i \in \{1, 2, \dots, M+1\}$. Substituting $i = M+1$ into (37) proves (35). The inequality (36) follows from (35) and the Lipschitz property (34). \square




Fig. 2. Results of algorithms and a locally polyhedral function

VI. EXPERIMENTS

In this section, Algorithm 1 (“Staggered”) is compared to the polynomial-decay averaging (“Polynomial”) in [5]. We also proposed another heuristic algorithm (“Heuristic”), which has a promising convergence rate. This heuristic algorithm replaces α in Algorithm 1 with

$$\alpha_t = \max[\alpha, c/(t+1)] \quad \forall t \in \{0, 1, 2, \dots\}$$

where c is some real-valued positive constant. This modification does not change the convergence rates in Sections III and IV, because it only adds $O(1/\alpha)$ rounds into the previous bounds.

For the purpose of comparison, Algorithm 1 uses $\alpha = 10^{-4}$. However, higher accuracy can be achieved by a smaller step size. The heuristic algorithm sets $c = 1$. The polynomial-decay averaging algorithm uses $c = 1$ and $\eta = 3$ (defined in [5]). Note that the stochastic subgradient algorithm with a constant step size (“Constant”) is the by product of Algorithm 1.

A locally polyhedral function $F = \|\mathbf{w}\|_1$ is considered where $\mathbf{w} \in [-4, 4]^{100}$. When $\mathbf{g}(\mathbf{w})$ is a subgradient of F at \mathbf{w} , a stochastic subgradient is $\hat{\mathbf{g}}(\mathbf{w}) = \mathbf{g}(\mathbf{w}) \times X$ where X is a uniform random variable from 0 to 2, so $\mathbb{E}[\hat{\mathbf{g}}(\mathbf{w})|\mathbf{w}] = \mathbf{g}(\mathbf{w})$. Ten experiments are performed, and the average values are sampled at $\{2^k - 2\}_{k=1}^{18}$ (one round before Algorithm 1 resets the averages). Results are shown in Figure 2. Both axes of Figure 2 are in a log scale.

The plots of Algorithm 1 and the polynomial-decay algorithm cross each other, because the former has faster convergence rate. The subgradient algorithm with constant step size stops improving due to the fixed value of the step size. However, Algorithm 1 keeps improving after the stop. This can be explained by (20) where the average helps reducing the last term on the right-hand-side. The plot of the heuristic algorithm shows its convergence.

A general convex function $F(\mathbf{w}) = \sum_{i=1}^{100} F_i(w^{(i)})$ is considered where, for $i \in \{1, \dots, 100\}$, $w^{(i)}$ is the i -component of \mathbf{w} and

$$F_i(w^{(i)}) = \begin{cases} -w^{(i)}, & \text{if } w^{(i)} < 0 \\ (w^{(i)})^2, & \text{if } w^{(i)} \geq 0. \end{cases}$$




Fig. 3. Results of algorithms and a general convex function




Fig. 4. Results of algorithms and the function (38)

The i -component of a stochastic subgradient is $\hat{g}_i(w) = g_i(w) + Y$, where $g_i(w)$ is the i -component of the true subgradient of F at w and Y is a uniform random variable between -1 and 1. Simulation uses the same parameters as the locally polyhedral case. The results in Figure 3 have a similar trend as in the locally polyhedral case except that the plot of the stochastic subgradient algorithm with constant step size crosses the plot of the polynomial decay averaging.

Then we consider a non-smooth convex function $F(w) = \sum_{i=1}^{100} F_i(w^{(i)})$ where

$$F_i(w^{(i)}) = \begin{cases} w^{(i)} - \frac{10^{-6}}{2}, & \text{if } w^{(i)} \geq \frac{10^{-6}}{2} \\ -w^{(i)} - \frac{10^{-6}}{2}, & \text{if } w^{(i)} \leq -\frac{10^{-6}}{2} \\ 0, & \text{otherwise.} \end{cases} \quad (38)$$

This function has uncountable minimizers. The stochastic subgradient is the component-wise addition of the true subgradient and the uniform random variable Y . Simulation results are shown in Figure 4. Comparing these results to the results in Figure 2 shows the same trend of convergence rates even though function F in (38) does not satisfy the uniqueness assumption.

VII. CONCLUSION

This paper considers stochastic non-smooth convex optimization. We propose the staggered time average algorithm

and prove its performance. When a function with a unique minimum satisfies the locally polyhedral structure, the algorithm has $O(1/T)$ convergence rate. For a general convex function with a unique minimum, we derive a class of bounds on the convergence rate of the algorithm. For a special case of deterministic problems with the locally polyhedral structure, an algorithm with $O(1/\epsilon^{1/M})$ convergence is proposed.

REFERENCES

- [1] L. Rosasco, E. D. Vito, A. Caponnetto, M. Piana, and A. Verri, "Are loss functions all the same?" *Neural Computation*, vol. 16, no. 5, pp. 1063–1076, May 2004.
- [2] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society, Series B*, vol. 58, pp. 267–288, 1994.
- [3] Y. Nesterov, *Introductory lectures on convex optimization: a basic course*. London: Kluwer Academic, 2004.
- [4] A. Rakhlin, O. Shamir, and K. Sridharan, "Making gradient descent optimal for strongly convex stochastic optimization," in *ICML*, 2012.
- [5] O. Shamir and T. Zhang, "Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes," in *ICML*, 2013.
- [6] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan, "Stochastic convex optimization," in *COLT*, 2009.
- [7] S. Boyd and A. Mutapcic, "Stochastic subgradient method," Stanford University, Tech. Rep., 2008.
- [8] H. Ouyang and A. Gray, "Stochastic smoothing for nonsmooth minimizations: Accelerating sgd by exploiting structure," in *ICML*, 2012.
- [9] E. Hazan and S. Kale, "Beyond the regret minimization barrier: Optimal algorithms for stochastic strongly-convex optimization," *Journal of Machine Learning Research*, vol. 15, pp. 2489–2512, 2014.
- [10] J. V. Burke and M. C. Ferris, "Weak sharp minima in mathematical programming," *SIAM Journal on Control and Optimization*, vol. 31, no. 5, pp. 1340–1359, 1993.
- [11] B. Polyak, *Introduction to Optimization*. Optimization Software, 1987.
- [12] D. Bertsekas, A. Nedić, and A. Ozdaglar, *Convex Analysis and Optimization*. Athena Scientific, 2003.
- [13] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *ICML*, 2003.
- [14] M. J. Neely, "Energy-aware wireless scheduling with near optimal backlog and convergence time tradeoffs," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, Apr. 2015, pp. 91–99.
- [15] L. Huang and M. J. Neely, "Delay reduction via lagrange multipliers in stochastic network optimization," *Automatic Control, IEEE Transactions on*, vol. 56, no. 4, Apr. 2011.
- [16] S. M. Ross, *Stochastic Processes (Wiley Series in Probability and Statistics)*, 2nd ed. Wiley, Feb. 1995.
- [17] P. Tseng, "On accelerated proximal gradient methods for convex-concave optimization," *submitted to SIAM Journal on Optimization*, 2008.